

UNITED STATES PATENT APPLICATION
OF
GUY L. STEELE, JR.
FOR
COMPARATOR UNIT FOR COMPARING
VALUES OF FLOATING POINT OPERANDS

INCORPORATION BY REFERENCE

[001] Related U. S. Patent Application Serial No. not yet assigned, filed on even date herewith in the name of Guy L. Steele Jr. and entitled "Floating Point System That Represents Status Flag Information Within A Floating Point Operand," assigned to the assignee of the present application, is hereby incorporated by reference.

FIELD OF THE INVENTION

[002] The invention relates generally to systems and methods for performing floating point operations, and more particularly to systems and methods for comparing a pair of floating point operands each of which contains encoded floating point status information.

BACKGROUND OF THE INVENTION

[003] Digital electronic devices, such as digital computers, calculators, and other devices, perform arithmetic calculations on values in integer, or "fixed point," format, in fractional, or "floating point" format, or both. IEEE Standard 754, (hereinafter "IEEE Std. 754" or "the Standard") published in 1985 by the Institute of Electrical and Electronic Engineers, and adopted by the American National Standards Institute (ANSI), defines several standard formats for expressing values in floating point format, and a number of aspects regarding behavior of computation in connection therewith. In accordance with IEEE Std. 754, a representation in floating point format comprises a plurality of binary digits, or "bits," having the structure:

$$s e_{msb} \dots e_{lsb} f_{msb} \dots f_{lsb}$$

where bit "s" is a sign bit indicating whether the entire value is positive or negative, bits "e_{msb}...e_{lsb}" comprise an exponent field represent the exponent "e" in unsigned binary biased

format, and bits " $f_{msb} \dots f_{lsb}$ " comprise a fraction field that represents the fractional portion "f" in unsigned binary format ("msb" represents "most significant bit" and "lsb" represents "least significant bit"). The Standard defines two general formats, namely, a "single" format which comprises thirty-two bits, and a "double" format which comprises sixty-four bits. In the single format, there is one sign bit "s," eight bits " $e_7 \dots e_0$ " comprising the exponent field, and twenty-three bits " $f_{22} \dots f_0$ " comprising the fraction field. In the double format, there is one sign bit "s," eleven bits " $e_{10} \dots e_0$ " comprising the exponent field and fifty-two bits " $f_{51} \dots f_0$ " comprising the fraction field.

[004] As indicated above, the exponent field of the floating point representation " $e_{msb} \dots e_{lsb}$ " represents the exponent "E" in biased format. The biased format provides a mechanism by which the sign of the exponent is implicitly indicated. In particular, the bits " $e_{msb} \dots e_{lsb}$ " represent a binary encoded value "e" such that " $e = E + \text{bias}$." This allows the exponent E to extend from -126 to +127, in the eight-bit "single" format, and from -1022 to +1023 in the eleven-bit "double" format, and provides for relatively easy manipulation of the exponents in multiplication and division operations, in which the exponents are added and subtracted, respectively.

[005] IEEE Std. 754 provides for several different formats with both the single and double formats, which are generally based on the bit patterns of the bits " $e_{msb} \dots e_{lsb}$ " comprising the exponent field and the bits " $f_{msb} \dots f_{lsb}$ " comprising the fraction field. For example, as shown in prior art FIG. 3, if a number is represented such that all of the bits " $e_{msb} \dots e_{lsb}$ " of the exponent field are binary ones (that is, if the bits represent a binary-encoded value of "255" in the single format or "2047" in the double format) and all of the bits " $f_{msb} \dots f_{lsb}$ " of the fraction field are binary zeros, then the value of the number is positive

infinity 310 or negative infinity 320, depending on the value of the sign bit "s"; in particular, the value "v" is $v = (-1)^s \infty$, where " ∞ " represents the value "infinity." On the other hand, if all of the bits " $e_{msb} \dots e_{lsb}$ " of the exponent field are binary ones and if the bits $f_{msb} \dots f_{lsb}$ of the fraction field are not all zeros, then the value that is represented is deemed "not a number," 330, abbreviated in the Standard by "NaN."

[006] If a number has an exponent field in which the bits " $e_{msb} \dots e_{lsb}$ " are neither all binary ones nor all binary zeros (that is, if the bits represent a binary-encoded value between 1 and 254 in the single format or between 1 and 2046 in the double format), the number is said to be in a "normalized" format 360. For a number in the normalized format, the value represented by the number is $v = (-1)^s 2^{e-bias} (1.f_{msb} \dots f_{lsb})$, where "|" represents a concatenation operation. Effectively, in the normalized format, there is an implicit most significant digit having the value "one," so that the twenty-three digits in the fraction field of the single format, or the fifty-two digits in the fraction field of the double format, will effectively represent a value having twenty-four digits or fifty-three digits of precision, respectively, where the value is less than two, but not less than one.

[007] On the other hand, if a number has an exponent field in which the bits " $e_{msb} \dots e_{lsb}$ " are all binary zeros, representing the binary-encoded value of "zero," and a fraction field in which the bits $f_{msb} \dots f_{lsb}$ are not all zero, the number is said to be in a "de-normalized" format 370. For a number in the de-normalized format, the value represented by the number is $v = (-1)^s 2^{e-bias+1} (0.f_{msb} \dots f_{lsb})$. It will be appreciated that the range of values of numbers that can be expressed in the de-normalized format is disjoint from the range of values of numbers that can be expressed in the normalized format, for both the single and double formats. Finally, if a number has an exponent field in which the bits " $e_{msb} \dots e_{lsb}$ " are

all binary zeros, representing the binary-encoded value of "zero," and a fraction field in which the bits $f_{msb} \dots f_{lsb}$ are all zero, the number has the value "zero." It will be appreciated that the value "zero" may be positive zero 340 or negative zero 350, depending on the value of the sign bit.

[008] Generally, floating point units to perform computations whose results conform to IEEE Std. 754 are designed to generate a result in response to a floating point instruction in three steps:

[009] (a) First, an approximation calculation step in which an approximation to the absolutely accurate mathematical result (assuming that the input operands represent the specific mathematical values as described by IEEE Std. 754) is calculated. This calculation is sufficiently precise as to allow this accurate mathematical result to be summarized by a sign bit, an exponent (typically represented using more bits than are used for an exponent in the standard floating-point format), and some number "N" of bits of the presumed result fraction, plus a guard bit and a sticky bit. The value of the exponent will be such that the value of the fraction generated in step (a) consists of a 1 before the binary point and a fraction after the binary point. The bits are calculated so as to obtain the same result as the following conceptual procedure (which is impossible under some circumstances to carry out in practice): calculate the mathematical result to an infinite number of bits of precision in binary scientific notation, and in such a way that there is no bit position in the significand such that all bits of lesser significance are 1-bits (this restriction avoids the ambiguity between, for example, 1.100000... and 1.011111... as representations of the value "one-and-one-half"); then let the N most significant bits of the infinite significand be used as the intermediate result significand, let the next bit of the infinite significand be the guard bit, and

let the sticky bit be 0 if and only if ALL remaining bits of the infinite significant are 0-bits (in other words, the sticky bit is the logical OR of all remaining bits of the infinite fraction after the guard bit).

[010] (b) Second, a rounding step, in which the guard bit, the sticky bit, perhaps the sign bit, and perhaps some of the bits of the presumed significand generated in step (a) are used to decide whether to alter the result of step (a). For the rounding modes defined by IEEE Std. 754, this is a decision as to whether to increase the magnitude of the number represented by the presumed exponent and fraction generated in step (a). Increasing the magnitude of the number is done by adding 1 to the significand in its least significant bit position, as if the significand were a binary integer. It will be appreciated that, if the significand is all 1-bits, then magnitude of the number is "increased" by changing it to a high-order 1-bit followed by all 0-bits and adding 1 to the exponent. It will be further appreciated that,

[011] (i) if the result is a positive number, and

[012] (a) if the decision is made to increase, effectively the decision has been made to increase the value of the result, thereby rounding the result up (that is, towards positive infinity), but

[013] (b) if the decision is made not to increase, effectively the decision has been made to decrease the value of the result, thereby rounding the result down (that is, towards negative infinity);
and

[014] (ii) if the result is a negative number, and

[015] (a) if the decision is made to increase, effectively the decision has been made to decrease the value of the result, thereby rounding the result down, but

[016] (b) if the decision is made not to increase, effectively the decision has been made to increase the value of the result, thereby rounding the result up.

[017] (c) Finally, a packaging step, in which the result is packaged into a standard floating-point format. This may involve substituting a special representation, such as the representation defined for infinity or NaN if an exceptional situation (such as overflow, underflow, or an invalid operation) was detected. Alternatively, this may involve removing the leading 1-bit (if any) of the fraction, because such leading 1-bits are implicit in the standard format. As another alternative, this may involve shifting the fraction in order to construct a denormalized number. As a specific example, we assume that this is the step that forces the result to be a NaN if any input operand is a NaN. In this step, the decision is also made as to whether the result should be an infinity. It will be appreciated that, if the result is to be a NaN or infinity, any result from step (b) will be discarded and instead the appropriate representation will be provided as the result.

[018] In addition in the packaging step, floating-point status information is generated, which is conventionally stored in a floating point status register. The floating point status information generated for a particular floating point operation includes indications, for example, as to whether:

[019] (i) a particular operand is invalid for the operation to be performed ("invalid operation");

[020] (ii) if the operation to be performed is division, the divisor is zero ("division-by-zero");

[021] (iii) an overflow occurred during the operation ("overflow");

[022] (iv) an underflow occurred during the operation ("underflow"); and

[023] (v) the rounded result of the operation is not exact ("inexact").

[024] These conditions are typically represented by flags that are stored in the floating point status register, separate from the floating point operand. The floating point status information can be used to dynamically control the operations in response to certain instructions, such as conditional branch, conditional move, and conditional trap instructions that may be in the instruction stream subsequent to the floating point instruction. Also, the floating point status information may enable processing of a trap sequence, which will interrupt the normal flow of program execution. In addition, the floating point status information may be used to affect certain ones of the functional unit control signals that control the rounding mode. IEEE Std. 754 also provides for accumulating floating point status information from, for example, results generated for a series or plurality of floating point operations.

[025] IEEE Std. 754 has brought relative harmony and stability to floating-point computation and architectural design of floating-point units. Moreover, its design was based on some important principles and rests on sensible mathematical semantics that ease the job of programmers and numerical analysts. It also supports the implementation of interval arithmetic, which may prove to be preferable to simple scalar arithmetic for many tasks. Nevertheless, IEEE Std. 754 has some serious drawbacks, including:

[026] (i) Modes, which include the rounding mode and may also include a traps enabled/disabled mode, flags representing the floating point status information that is stored in the floating point status register, and traps that are required to implement IEEE Std. 754 all introduce implicit serialization between floating-point instructions, and between floating point instructions and the instructions that read and write the flags and modes. Rounding modes can introduce implicit serialization because they are typically indicated as a global state, although in some microprocessor architectures, the rounding mode is encoded as part of the instruction operation code, which will alleviate this problem to some extent. This implicit serialization makes the Standard difficult to implement coherently in today's superscalar and parallel microprocessor architectures without loss of performance.

1. Implicit serialization occurs when programmers and designers try to avoid the problems caused if every floating point instructions uses, and can change, the same floating point status register. This can create problems if, for example, two instructions are executing in parallel in a microprocessor architectures featuring several CPUs running at once and both cause an update of the floating point status register. In such a case, the contents of the status register would likely be incorrect with respect to at least one of the instructions, because the other parallel instruction will have written over the original contents. Similar problems can occur in scalar processor architectures, in which several instructions are issued and processed at once. To solve this problem, programmers and designers make sure that floating point instructions that can affect the floating point status register are executed in serial fashion, one instruction completing before another begins.

[027] (ii) The implicit side effects of a procedure that can change the flags or modes can make it very difficult for compilers to perform optimizations on floating-point

code. To be safe, compilers for most languages assume that every procedure call is an optimization barrier.

[028] (iii) Global flags, such as those that signal certain modes, make it more difficult to do instruction scheduling where the best performance is provided by interleaving instructions of unrelated computations. Instructions from regions of code governed by different flag settings or different flag detection requirements cannot easily be interleaved when they must share a single set of global flag bits.

[029] (iv) Traps have been difficult to integrate efficiently into architectures and programming language designs for fine-grained control of algorithmic behavior.

[030] U. S. Patent Application Serial No. not yet assigned, filed on even date herewith in the name of Guy L. Steele Jr. and entitled "Floating Point System That Represents Status Flag Information Within A Floating Point Operand," describes a floating point unit in which floating point status information is encoded in the representations of the results generated thereby. By encoding the floating point status information relating to a floating point operation in the result that is generated for the operation, the implicit serialization required by maintaining the floating point status information separate and apart therefrom can be obviated. In one embodiment, the floating point unit includes a plurality of functional units, including an adder unit, a multiplier unit, a divider unit, a square root unit, a maximum/minimum unit, a comparator unit and a tester unit, all of which operate under control of functional unit control signals provided by a control unit. Systems and methods consistent with the principles of the present invention may provide a comparator unit that can be used in the floating point unit described in that application.

SUMMARY OF THE INVENTION

[031] Systems and methods consistent with the principles of the invention provides new and improved ways of performing floating point operations, and more particularly, comparing two floating point numbers by selectively providing a signal that indicates the comparative relationship among floating point operands, taking into consideration floating point status data encoded in the operands and/or the overall format of operands.

In one embodiment consistent with the principles of the present invention, a floating point comparator circuit for comparing a plurality of floating point operands includes a plurality of analysis circuits, one for each of the floating point operands, configured to determine a format of each of the floating point operands based upon floating point status information encoded within each of the floating point operands, and a result generator circuit coupled to the analysis circuits, the result generator circuit configured to generate a result signal based on the format determined by each analysis circuit and based on a comparative relationship among the floating point operands. The format of each of the floating point operands may be from a group comprising: not-a-number (NaN), infinity, normalized, denormalized, zero, invalid operation, overflow, underflow, division by zero, exact, and inexact. The result generator circuit may ignore the encoded floating point statuses of the plurality of floating point operands when comparing just the magnitudes of the plurality of floating point operands.

BRIEF DESCRIPTION OF THE DRAWINGS

[032] This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

[033] FIG. 1 is a functional block diagram of an exemplary comparator unit constructed in accordance with an embodiment of the invention;

[034] FIG. 2 depicts formats for exemplary representations of floating point values used by the comparator unit depicted in FIG. 1 and in accordance with an embodiment of the invention; and

[035] FIG. 3 depicts prior art formats for representation of floating point values.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

[036] FIG. 1 is a functional block diagram of an exemplary comparator unit 10 constructed in accordance with an embodiment of the invention. Generally, the comparator unit 10 receives two floating point operands and, in response during a comparison operation, generates signals that indicate whether the value of one operand is greater than the value of the other, the value of one operand is less than the value of the other, the values of the two operands are equal, or the values of the two operands cannot be compared. As will be described below, in some cases, operands of comparator unit 10 can include floating point status information, with the floating point status information being encoded in and comprising part of the floating point representation of the result. Since the floating point status information comprises part of the operand, instead of being separate and apart from the result as in prior art comparator units, the implicit serialization that is required by maintaining the floating point status information separate and apart from the result can be obviated. The comparator unit 10 facilitates the comparison of values of the respective operands in a manner so as to accommodate the floating point status information that may be included therein.

[037] The comparator circuit 10 uses the encoded floating point status information in results that are generated in certain formats. This will be illustrated in connection with FIG. 2. FIG. 2 depicts formats of floating point operands that the comparator circuit 10 may receive. With reference to FIG. 2, seven exemplary formats are depicted, including a zero format 100, an underflow format 101, a denormalized format 102, a normalized non-zero format 103, an overflow format 104, an infinity format 105 and a not-a-number (NaN) format 106. The zero format 100 is used to represent the values "zero," or, more specifically, positive or negative zero, depending on the value of "s," the sign bit.

[038] The exemplary underflow format 101 provides a mechanism by which the floating point comparator circuit 10 can indicate that the result of a computation is an underflow. In the underflow format of the illustrated embodiment, the sign bit "s" indicates whether the result is positive or negative, the bits $e_{msb} \dots e_{lsb}$ of the exponent field are all binary zeros, and the bits $f_{msb} \dots f_{lsb+1}$ of the fraction field are all binary zeros. The least significant bit f_{lsb} of the fraction field is a binary one.

[039] The exemplary denormalized format 102 and exemplary normalized non-zero format 103 are used to represent finite non-zero floating point values substantially along the lines of that described above in connection with IEEE Std. 754. In both formats 102 and 103 of the illustrated embodiment, the sign bit "s" indicates whether the result is positive or negative. The bits $e_{msb} \dots e_{lsb}$ of the exponent field of the denormalized format 102 are all binary zeros, whereas the bits $e_{msb} \dots e_{lsb}$ of the exponent field of the normalized non-zero format 103 are mixed ones and zeros, except that the exponent field of the normalized non-zero format 103 will not have a pattern in which bits $e_{msb} \dots e_{lsb+1}$ are all binary ones and the

least significant bit e_{lsb} zero and all of the bits $f_{msb} \dots f_{lsb}$ of the fraction field are all binary ones. In format 102 the bits $f_{msb} \dots f_{lsb}$ of the fraction field are not all binary zeros.

[040] The exemplary overflow format 104 provides a mechanism by which the floating point comparator circuit 10 can indicate that the result of a computation is an overflow. In the overflow format 104 of the illustrated embodiment, the sign bit "s" indicates whether the result is positive or negative, the bits $e_{msb} \dots e_{lsb+1}$ of the exponent field are all binary ones, with the least significant bit e_{lsb} being zero. The bits $f_{msb} \dots f_{lsb}$ of the fraction field are all binary ones.

[041] The exemplary infinity format 105 provides a mechanism by which the floating point comparator circuit 10 can indicate that a result is infinite. In the infinity format 105 of the illustrated embodiment, the sign bit "s" indicates whether the result is positive or negative, the bits $e_{msb} \dots e_{lsb}$ of the exponent field are all binary ones, and the bits $f_{msb} \dots f_{lsb}$ of the fraction field are all binary zero's. The five least significant bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field are flags, which will be described below.

[042] The exemplary NaN (not-a-number) format 105 provides a mechanism by which the floating point comparator circuit 10 can indicate that the result is not a number. In the NaN format of the illustrated embodiment, the sign bit "s" can be any value, the bits $e_{msb} \dots e_{lsb}$ of the exponent field are all binary ones, and the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field are not all binary zeros. The five least significant bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field are flags, which will be described below.

[043] As noted above, in values represented in the infinity format 105 and the NaN format 106, the five low order bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field are flags. In the formats used with the floating point comparator circuit 10, the five flags include the flags that are

defined by IEEE Std. 754, including an invalid operation flag "n," an overflow flag "o," an underflow flag "u," a division-by-zero flag "z," and an inexact flag "x." For example, a value in the NaN format 106 in which both the overflow flag "o" and the division-by-zero flag "z" are set indicates that the value represents a result of a computation that involved an overflow (this from the overflow flag "o"), as well as an attempt to divide by zero (this from the division-by-zero flag "z"). It should be noted that the flags provide the same status information as provided by, for example, information stored in a floating point status register in a prior art floating point unit. However, since the information is provided as part of the result and stored therewith in registers in which the result is ultimately stored, multiple instructions can be contemporaneously executed. This may be accomplished because the floating point status information generated during execution of one instruction, when stored, will not over-write previously-stored floating point status information generated during execution of another instruction. In addition, in one embodiment, values in the other formats can be indicated as being inexact based in part on the least-significant bit f_{lsb} of their fraction fields, and in that embodiment that bit operates as an inexact flag. The value will be indicated as being inexact if the bit f_{lsb} has the value "one," and otherwise exact.

[044] One of ordinary skill in the art will recognize that the other formats and bit patterns could be used to represent the floating point operand formats without departing from the principles of the present invention. One of ordinary skill in the art will also recognize that the floating point status information contained in the operands could easily be represented by other bit combinations (not shown) without departing from the principles of the present invention. For example, more or fewer bits could be used, a subset or superset of the exemplary status bits could be used, or the most significant bits of an operand (or some

other subset of bits) could be used to indicate the floating point status information, instead of the least significant bits illustrated.

[045] With this background the structure and operation of the exemplary comparator unit 10 will be described in connection with FIG. 1. With reference to FIG. 1, the comparator unit 10 includes two operand buffers 11A and 11B, respective operand analysis circuits 12A and 12B, and a result generator 13. The operand buffers 11A and 11B receive and store respective operands that may be received from, for example, a set of registers (not shown) in a conventional manner. Each operand analysis circuit 12A, 12B analyzes the operand in the respective operand buffer 11A, 11B and generates signals providing information relating to the respective operands. For decision circuit 13, these signals essentially provide information as to the type of the respective operand. In particular, these signals may indicate whether the respective operand is in the zero format 100, the infinity format 105, or the NaN format 106. The result generator 13 receives the signals from the operand analysis circuits 12A, 12B and additional signals representative of the operands in the operand buffers 11A and 11B and generates four result signals. In more detail, the four result signals include an A LT B (operand A is less than operand B) signal, an A GT B (operand A is greater than operand B) signal, an A EQ B (operand A is equal to operand B) signal, and an INCOMPARABLE (operand A cannot be compared to operand B) signal. One of these four signals will be asserted to indicate the comparative relationship between the values of the operands in the operand buffers 11A and 11B. If the INCOMPARABLE signal is asserted, operand A cannot be compared to operand B. Operand A cannot be compared to operand B if, for example, one or both of the operands is in the NaN format 106. In one embodiment, if the operands in both operand buffers 11A and 11B are in the zero

format 100, they will be equal even if their signs differ. The comparator unit 10 can provide the result signals to a control unit (not shown) for controlling a floating point unit that includes the comparator unit 10, as described in the aforementioned related patent application.

[046] One of ordinary skill in the art will recognize that the comparative relationship between two operands can be easily redefined to differ from those of the exemplary embodiments, without departing from the scope of the present invention. For example, the aforementioned “equal” comparative relationship between positive zero and negative zero may be redefined such that positive zero is “greater than” negative zero. Or, for example, the comparative relationship between a finite value operand and a NaN format operand may be defined such that a finite value is greater than NaN (instead of incomparable), without departing from the scope of the present invention.

[047] As noted above, each operand analysis circuit 12A, 12B analyzes the operand in the respective buffer 11A, 11B and generates signals providing information relating thereto, which signals are provided to the decision circuit 13. In the embodiment illustrated in FIG. 1, each exemplary operand analysis circuit 12A, 12B, generally comprises a number of comparators, including:

[048] (i) a comparator 20A, 20B that generates an asserted signal if the bits $e_{msb} \dots e_{lsb}$ of the exponent field of the operand in respective buffer 11A, 11B are all binary ones, which will be the case if the operand is in the infinity format 105 or the NaN format 106;

[049] (ii) a comparator 21A, 21B that generates an asserted signal if the bits $e_{msb} \dots e_{lsb}$ of the exponent field of the operand in respective buffer 11A, 11B are all binary

zeros, which will be the case if the operand is in the zero format 100, underflow format 101, or denormalized format 102;

[050] (iii) a comparator 22A, 22B that generates an asserted signal if the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in the respective buffer 11A, 11B are all binary zeros, which may be the case if the operand is in the zero format 100, underflow format 101, denormalized format 102, normalized non-zero format 103 or infinity format 105;

[051] (iv) a comparator 23A, 23B that generates an asserted signal if all of the bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in the respective buffer 11A, 11B are binary zero's, which will be the case if the operand is in the zero format 100, and which may be the case if the operand is in the denormalized format 102, normalized non-zero format 103, or if the flags "n," "o," "u," "z" and "x" are clear in the infinity format 105 or NaN format 106;

[052] (v) a NAND gate 24A, 24B that generates an asserted A NaN, B NaN signal if the comparator 20A, 20B is generating an asserted signal and the comparator 22A, 22B is generating a negated signal; it will be appreciated that the NAND gate 24A, 24B will be generating an asserted signal if the operand in the respective operand buffer 11A, 11B is in the NaN format 106;

[053] (vi) an AND gate 25A, 25B that generates an asserted A INF, B INF signal if the comparator 22A, 22B is generating an asserted signal and the comparator 20A, 20B is also generating an asserted signal; it will be appreciated that the AND gate 25A, 25B will be generating an asserted signal if the operand in the respective operand buffer 11A, 11B is in the infinity format 105; and

[054] (vii) an AND gate 26A, 26B that generates an asserted A ZERO, B ZERO signal if the comparator 21A, 21B is generating an asserted signal, the comparator 22A, 22B

is generating an asserted signal, and the comparator 23A, 23B is generating an asserted signal; it will be appreciated that this will be the case if both of the operands in the operand buffers 11A, 11B are in the zero format 100.

[055] As noted above, the result generator 13 receives the signals from the operand analysis circuits 12A, 12B and signals representative of the operands in the operand buffers 11A and 11B and generates four result signals. In one embodiment, these result signals may include an A LT B (operand A is less than operand B) signal, an A GT B (operand A is greater than operand B) signal, an A EQ B (operand A is equal to operand B) signal, and an INCOMPARABLE signal, one of which will be asserted to indicate the comparative relationship between the values of the operands in the operand buffers 11A and 11B.

[056] In the illustrated embodiment of the invention in FIG. 1, the exemplary result generator 13 includes:

[057] (viii) an OR gate 30 that generates an asserted INCOMPARABLE signal if either the A NaN signal from NAND gate 24A or the B NaN signal from NAND gate 24B is asserted; it will be appreciated that the OR gate 30 will be generating an asserted signal if the operand in either operand buffer 11A, 11B is in the NaN format 106;

[058] (ix) a comparator 31 that generates an asserted B GT A (operand B is greater than operand A) signal if the magnitude, or absolute value, of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A;

[059] (x) a comparator 32 that generates an asserted B EQ A (operand B is equal to operand A) signal if the magnitude of the operand in the operand buffer 11 B is equal to the magnitude of the operand in operand buffer 11A;

[060] (xi) an XOR gate 34 that generates an asserted S A NEQ S B (sign A not equal to sign B) signal if an A NEG signal, representative of the sign bit of the operand in operand buffer 11A, is asserted and a B NEG signal, representative of the sign bit of the operand in operand buffer 11B, is negated, or vice versa; otherwise XOR gate 34 generates a negated signal; it will be appreciated that the XOR gate 34 will be generating an asserted signal if the operands in the operand buffers 11A and 11B are of opposite signs, and a negated signal if the operands are of the same sign;

[061] (xii) a NAND gate 35 that generates an asserted signal if the A INF signal from AND gate 25A and the B INF signal from AND gate 25B are both asserted and the S A NEQ S B (sign A not equal to sign B) signal from XOR gate 34 is negated; it will be appreciated that the NAND gate 35 will be generating an asserted signal if the operands in both operand buffers 11A and 11B are in the infinity format 105 and are of the same sign;

[062] (xii) a NAND gate 36 that generates an asserted signal if the INCOMPARABLE signal from OR gate 30 is negated and the B EQ A signal from comparator 32 is asserted and the S A NEQ S B (sign A not equal to sign B) signal from XOR gate 34 is negated; it will be appreciated that the NAND gate 36 will be generating an asserted signal if the magnitude and sign of the operand in operand buffer 11A equals the magnitude and sign of the operand in operand buffer 11B and neither operand is in the NaN format 106;

[063] (xiv) an AND gate 37 that generates an asserted signal if the A ZERO and B ZERO signals from AND gates 26A and 26B are both asserted; it will be appreciated that the AND gate 37 will be generating an asserted signal if the operands in both operand buffers 11A and 11B are in the zero format 100, regardless of the signs of the respective operands;

[064] (xv) an OR gate 38 that generates an asserted A EQ B (operand A equals operand B) signal if any of NAND gates 35 or 36 or AND gate 37 is generating an asserted signal; it will be appreciated that OR gate 38 will be generating an asserted signal if:

[065] (a) the operands in both operand buffers 11A and 11B are in the infinity format 105 and are of the same sign (reference NAND gate 35); or

[066] (b) the magnitudes and signs of the operands in the operand buffers 11A and 11B are the same and neither operand is in the NaN format 106 (reference NAND gate 36); or

[067] (c) the operands in both operand buffers 11A and 11B are in the zero format 100, regardless of their signs (reference AND gate 37);

[068] (xvi) a NAND gate 39 that generates an asserted signal if the A NEG signal representative of the sign of the operand in operand buffer 11A is asserted and the B NEG signal representative of the sign of the operand in operand buffer 11B is not asserted; it will be appreciated that the NAND gate 39 will be generating an asserted signal if the sign of the operand in operand buffer 11A is negative and the sign of the operand in operand buffer 11B is positive;

[069] (xvii) an XOR gate 40 that generates an asserted signal if the B GT A signal from the comparator 31 is asserted and the B NEG signal is negated, or if the B GT A signal is negated and the B NEG signal is asserted, it will be appreciated that the XOR gate 40 will be generating:

[070] (a) an asserted signal if

[071] (I) the magnitude of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A and the operand in operand buffer 11B is positive, or

[072] (II) the magnitude of the operand in operand buffer 11B is not greater than the magnitude of the operand in operand buffer 11A and the operand in operand buffer 11B is negative, and

[073] (b) a negated signal if

[074] (1) the magnitude of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A and the operand in operand buffer 11B is negative, or

[075] (II) the magnitude of the operand in operand buffer 11B is not greater than the magnitude of the operand in operand buffer 11A and the operand in operand buffer 11B is positive;

[076] (xviii) a NAND gate 41 that generates an asserted signal if the signal generated by the XOR gate 34 is negated and the signal generated by XOR gate 40 is asserted; it will be appreciated that NAND gate 41 will be generating an asserted signal if:

[077] (a) the A NEG and B NEG signals are both negated indicating that the operands in both operand buffers are positive, and the magnitude of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A, or

[078] (b) the A NEG and B NEG signals are both asserted, indicating that the operands in both operand buffers are negative, and the

magnitude of the operand in operand buffer 11B is not greater than the magnitude of the operand in operand buffer 11A;

[079] (xix) an OR gate 42 that generates an asserted signal if either NAND gate 39 or NAND gate 41 is generating an asserted signal, it will be appreciated that OR gate 42 will be generating an asserted signal if:

[080] (a) the A NEG and B NEG signals are both negated, indicating that the operands in both operand buffers are positive, and the magnitude of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A, or

[081] (b) the A NEG and B NEG signals are both asserted, indicating that the operands in both operand buffers are negative, and the magnitude of the operand in operand buffer 11B is not greater than the magnitude of the operand in operand buffer 11A, or

[082] (c) the operand in buffer 11A is negative and the operand in buffer 11B is positive;

[083] (xx) a NAND gate 43 that generates an asserted A LT B operand A less than operand B signal if OR gate 42 is generating an asserted signal and OR gates 30 and 38 are generating negated signals; it will be appreciated that NAND gate 43 will be generating an asserted signal if the operands can be compared, the operands are not equal; and:

[084] (a) the A NEG and B NEG signals are both negated, indicating that the operands in both operand buffers are positive, and the magnitude of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A, or

- [085] (b) the A NEG and B NEG signals are both asserted, indicating that the operands in both operand buffers are negative, and the magnitude of the operand in operand buffer 11B is not greater than the magnitude of the operand in operand buffer 11A, or
- [086] (c) the operand in buffer 11A is negative and the operand in buffer 11B is positive;

[087] (xxi) a NAND gate 44 that generates an asserted A GT B (operand A greater than operand B) signal if OR gates 42, 30 and 38 are generating negated signals; it will be appreciated that NAND gate 44 will be generating an asserted signal if the operands can be compared and the operands are not equal, and either:

- [088] (a) the operands in both operand buffers are positive, and the magnitude of the operand in operand buffer 11B is not greater than the magnitude of the operand in operand buffer 11A, or
- [089] (b) the operands in both operand buffers are negative, and the magnitude of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A, or
- [090] (c) the operand in buffer 11A is positive or the operand in buffer 11B is negative.

[091] As noted above, the exemplary comparator unit 10 can provide the A LT B, A GT B, A EQ B and INCOMPARABLE result signals to a control unit for controlling a floating point unit that includes the comparator unit 10, as described in the aforementioned related patent application. The result of a comparison operation can be used for a number of functions, which will be apparent to those skilled in the art, including controlling the

behavior of conditional branch, conditional move, boolean result generation, and conditional trap instructions. The instruction that enables the comparison operation to be performed may be a separate instruction, or the comparison operation may be performed as part of the execution of a conditional branch, conditional move, boolean result generation or conditional trap instruction.

[092] One of ordinary skill in the art will recognize that the circuits implementing the comparative relationships between the operands specified in the preceding description to produce the result can easily be altered to implement different comparative relationships and generate different results without departing from the principles of the present invention. For example, the relationships may be redesigned so that the generated result is A LT B if operand A is in the negative zero format and operand B is in the positive zero format, instead of generating an A EQ B result as is done by the exemplary circuit.

[093] It will be appreciated that a system in accordance with the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof, any portion of which may be controlled by a suitable program. Any program may in whole or in part comprise part of or be stored on the system in a conventional manner, or it may in whole or in part be provided to the system over a network or other mechanism for transferring information in a conventional manner. In addition, it will be appreciated that the system may be operated and/or otherwise controlled by means of information provided by an operator using operator input elements (not shown) which may be connected directly to the system or which may transfer the information to the system over a network or other mechanism for transferring information in a conventional manner.

[094] Those skilled in the art will appreciate that the invention may be practiced in an electrical circuit comprising discrete electronic elements, packaged or integrated electronic chips containing logic gates, a circuit utilizing a microprocessor, or on a single chip containing electronic elements or microprocessors. It may also be provided using other technologies capable of performing logical operations such as, for example, AND, OR, and NOT, including but not limited to mechanical, optical, fluidic, and quantum technologies. In addition, the invention may be practiced within a general purpose computer or in any other circuits or systems as are known by those skilled in the art.

[095] The foregoing description has been limited to specific embodiments of this invention. It will be apparent, however, that various variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention. It is the object of the appended claims to cover these and such other variations and modifications as come within the true spirit and scope of the invention